
Subject: Re: [RFC][PATCH 14/16] Introduce proc_mnt for pid_ns
Posted by [Pavel Emelianov](#) on Fri, 25 May 2007 12:27:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dave Hansen wrote:

> On Thu, 2007-05-24 at 14:15 +0400, Pavel Emelianov wrote:

>>> s->s_flags |= MS_NODIRATIME | MS_NOSUID | MS_NOEXEC;

>>> @@ -466,6 +467,7 @@ int proc_fill_super(struct super_block *

>>> s->s_magic = PROC_SUPER_MAGIC;

>>> s->s_op = &proc_sops;

>>> s->s_time_gran = 1;

>>> + s->s_fs_info = pid_ns;

>> One more thing I've just noticed - you don't get the namespace

>> here so after all the tasks die and namespace is freed we

>> have a proc mount pointing to freed namespace...

>

> Yep, missed reference count. Thanks for catching this!

refcount is not the only badness. As I said you have
race in testing superblock's namespace and setting it
and one more: do you use one struct proc_dir_entry
proc_root for all the superlocks?

> -- Dave

>

>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
